

CLAIMS

What is claimed is:

- 5        1. An asymmetric data processor comprising:
  - a first group of nodes comprising one or more host processors, each host comprising a memory, a network interface, and one or more Central Processing Units (CPUs), wherein each host accepts and responds to queries for data, and transforms such queries into one or more jobs;
  - 10        a second group of nodes comprising one or more Job Processing Units (JPUs), wherein each JPU comprises:
    - a memory, for storing data
    - a network interface, for receiving data and instructions
    - a streaming data interface, for receiving data from a streaming data source;
    - 15        one or more general purpose CPUs, for responding to requests from at least one host computer in the first group, and to requests from other JPUs in the second group, and
    - 20        one or more Programmable Streaming Data Processors (PSDPs), which perform primitive functions directly on data received from the streaming data interface, each PSDP thus performing initial processing on a set of data; and
    - a network connecting the nodes within each group and between the two groups, and
    - 25        wherein a JPU receives jobs from one or most nodes in the first group, performs work requested by the job, and forms a reply.
2. The apparatus of claim 1 wherein the data comprises structured records, and the structured records further comprise fields of various lengths and data types.

3. An apparatus as in claim 1 wherein the primitive functions performed by the PSDPs comprise field-level filtering.
4. An apparatus as in claim 1 wherein the streaming data interface is an industry-standard mass storage interface.  
5
5. An apparatus as in claim 2 in which at least one selected PSDP performs Boolean comparisons of record field values against other values.
- 10 6. An apparatus as in claim 5 wherein the Boolean comparision is against other record field values, and/or values held internally to that PSDP.
- 15 7. An apparatus as in claim 5 in which the selected PSDP restricts records that fail Boolean comparisons of field values, as such records stream into the PSDP and before such records are placed into the memory of the associated JPU.
- 20 8. An apparatus as in claim 2 in which the selected PSDP filters out fields of records that are not needed for particular queries, as such fields stream into the PSDP and before such fields are placed into the memory of the associated JPU, projecting forward into JPU memory those fields that are needed.
9. An apparatus as in claim 2 in which the PSDP output data may contain projected fields not contained in the source data, such as row address, transforms, results of expression evaluation, results of bit joins, and results of visibility tests.  
25
10. An apparatus as in claim 2 in which a selected PSDP decompresses fields and/or records.

11. An apparatus as in claim 1 wherein the streaming data interface is connected to receive data from a peripheral device selected from the group consisting of disk drive, network interface, and other streaming data source.
- 5 12. An apparatus as in claim 2 in which a selected PSDP performs a join operation, where the field values being joined have a small range of values, so that the presence or absence of a particular value can then be encoded as a bit within a sequence of bits, whose position within the sequence corresponds to the field value.
- 10 13. An apparatus as in claim 2 in which a selected PSDP performs an "exist join" operation, where the field values being joined have a small range of values, so that the presence or absence of a particular value can then be encoded as a bit within a sequence of bits, whose position within the sequence corresponds to the field value.
- 15 14. An apparatus as in claim 1 in which space is reserved in JPU memory at the head of the first tuple produced by the PSDP for recording tuple length and null vector, so that the length and null vectors from the end of the tuple may be relocated to this space.
- 20 15. An apparatus as in claim 1 in which at least one PSDP is implemented as a Field Programmable Gate Array (FPGA).
- 25 16. An apparatus as in claim 1 in which the host computers in the first group contain software comprising a plan optimizer component that determines which filtering primitives should be executed within a PSDP.
17. An apparatus as in claim 1 in which the JPUs in the second group contain software comprising a plan optimizer component that determines which filtering primitives should be executed within a PSDP.

18. An apparatus as in claim 1 in which the host computers in the first group contain software comprising a plan link component, which determines a query execution plan, the query execution plan further having portions that will be processed by a PSDP, 5 portions that will be processed by a JPU after a PSDP has returned data to the JPU, and portions that will be processed by a host, after the JPU has returned data to the host group.
19. An apparatus as in claim 1 in which the JPUs in the second group contain 10 software comprising a plan link component, which determines a query execution plan, the query execution plan further having portions that will be processed by a PSDP, portions that will be processed by a JPU after a PSDP has returned data to the JPU, and portions that will be processed by a host, after the JPU has returned data to the host group.
20. An apparatus as in claim 1 in which the hosts in the first group contain software 15 comprising a PSDPPrep component, which, for a given query execution plan, defines primitive instructions.
21. An apparatus as in claim 1 in which the JPUs in the second group contain 20 software comprising a PSDPPrep component, which, for a given query execution plan, defines primitive instructions.
22. An apparatus as in claim 21 wherein the instructions defined by the PSDPPrep 25 component include instructions to process fields of records.
23. An apparatus as in claim 21 in which a PSDPPrep component further identifies filtering, transformation, projection and/or aggregation operations to be performed by a PSDP.

24. An apparatus as in claim 21 in which a PSDPPrep component further modifies the query execution plan to specify restrict operations that are to be performed by a PSDP instead of a JPU.

5

25. An apparatus as in claim 1 in which the JPUs contain software comprising a PSDP Filter component, which loads an executable code image into a PSDP.

10

26. An apparatus as in claim 1 in which the JPUs contain software comprising a PSDP Scheduler component, which schedules jobs to run on a PSDP and queues PSDP requests to retrieve required data.

15

27. An apparatus as in claim 1 in which the JPUs in the second group contain software comprising a JPU Resource Scheduler component, which is responsible for scheduling jobs to be run on the JPU.

20

28. An apparatus as in claim 27 in which the JPU Resource Scheduler component further schedules jobs to run on a PSDP, communicating with a PSDP Scheduler component to queue up PSDP requests to retrieve required data.

25

29. An apparatus as in claim 27 in which the JPU Resource Scheduler component further schedules jobs, in which similar PSDP instructions in different query execution plans are combined to avoid duplicate PSDP processing requests.

30. An apparatus as in claim 2 in which an initial query is provided by a structured query language (SQL) statement, and the records specified thereby exist in various processing states within at least two components of the system including at least within a PSDP within a JPU, and/or within a host.

31. An apparatus as in claim 30 in which a PSDP processes fields within records are received from the streaming data source, without waiting to process any records until all records are received.